



1 This application is submitted in the name of the following inventor(s):

2

3 <u>Inventor</u>	4 <u>Citizenship</u>	5 <u>Residence City and State</u>
6 Banga, Gaurav	7 India	8 San Carlos, California

JC542 U.S. PTO
09/456027

9

10 The assignee is Network Appliance, Inc., a corporation having an office at

11 495 East Java Drive, Sunnyvale California 94089.

12 Title of the Invention

13 Computer Assisted Automatic Error Detection and Diagnosis of File Servers

14 Background of the Invention

15 1. *Field of the Invention*

16

17 This invention relates to computer assisted automatic error detection and

18 diagnosis of file servers, such as for a networked file server.

19

1 2. *Related Art*

2

3 Network file servers are subject to errors and other failures (such as less-
4 than-expected levels of performance), including those arising from hardware failure,
5 software error, erroneous configuration, or mismatch between configuration and usage.
6 Many of these errors and other failures are similar to those that are common for general-
7 purpose computer systems, and these would therefore be subject to similar forms of error
8 detection and diagnosis by a user of the system. However, many network file servers are
9 designed for minimal user intervention, and are specifically intended to have extremely
10 simple user interfaces. One such group of easy-to-use network file servers are called
11 “Network Appliance”™ file servers, for which the model of use is that the device is as
12 easy to use as a common toaster.

EU
GB
CA
13

14 One problem with known systems is that relative ease of use is often cou-
15 pled with relatively unsophisticated users. Error detection and diagnosis by relatively un-
16 sophisticated users is subject to the drawback that relatively unsophisticated users are
17 relatively unsophisticated in the use of error detection and diagnosis techniques. Thus,
18 forms of error detection and diagnosis that are acceptable for general-purpose computer
19 systems are not nearly as suitable for network file servers where one of the most impor-
20 tant purposes is to be easy-to-use.

21

1 Accordingly, it would be advantageous to provide a technique for computer
2 assisted automatic error detection and diagnosis of file servers that is not subject to draw-
3 backs of the known art.

4

5 Summary of the Invention

6

7 The invention provides a method and system for computer assisted auto-
8 matic error detection and diagnosis of file servers. In a preferred embodiment, the file
9 server includes diagnostic software modules for (1) periodic and continuous interpretation
10 of monitoring statistics, (2) augmentation of known network protocols, (3) cross-layer
11 analysis of monitoring statistics, and (4) tracking of hardware and software configuration
12 changes.

13

14 In a first aspect of the invention, the diagnostic software modules periodi-
15 cally and continuously review monitoring statistics gathered by the file server regarding
16 its operation. These monitoring statistics can include a wide variety of values gathered by
17 software modules at disparate levels within the operating system of the file server. The
18 collection of monitoring statistics is processed by a pattern matching system and a rule-
19 based inference system. In a preferred embodiment, the pattern matching system and
20 rule-based inference system are responsive to a use profile for the file server, such as
21 whether the file server is being used for an ISP, a development environment, a mail
22 server, or otherwise.

1
2 In a second aspect of the invention, the diagnostic software modules are ca-
3 pable of augmenting known network protocols, by manipulating parameters of lower-
4 level protocols using different higher-level protocols. The diagnostic software modules
5 can manipulate known parameters of the lower-level protocols in rapid succession, so as
6 to try a large number of combinations of protocol parameters. Using the higher-level
7 protocols, the diagnostic software modules can determine if the selected parameters for
8 the lower-level protocols are correct.

9
10 In a third aspect of the invention, the diagnostic software modules are capa-
11 ble of imposing sequential and combined constraints on diagnosis of possible errors, with
12 reference to known logical coupling between monitoring statistics gathered at multiple
13 logical levels of software modules within the file server. In a preferred embodiment, con-
14 straints from multiple logical levels are chained together so as to limit the number of pos-
15 sible errors deduced as possible from the various monitoring statistics to a relatively small
16 number.

17
18 In a fourth aspect of the invention, the diagnostic software modules are ca-
19 pable of tracking hardware and software configuration changes to the file server, and re-
20 lating changes in known monitoring statistics to timing of those hardware and software
21 configuration changes. In a preferred embodiment, the diagnostic software modules are
22 capable of determining the configuration change most likely to be responsible for a com-

1 computer assisted diagnosed error, and of suggesting activities to reverse the hardware and
2 software configuration changes so as to place the file server in an operating state.

3

4 The invention provides an enabling technology for a wide variety of appli-
5 cations for computer assisted automatic error detection and diagnosis of file servers, so as
6 to obtain substantial advantages and capabilities that are novel and non-obvious in view
7 of the known art. Examples described below primarily relate to networked file servers,
8 but the invention is broadly applicable to many different types of automated software
9 systems.

10

11 Brief Description of the Drawings

12

13 Figure 1 shows a block diagram of a system for computer assisted auto-
14 matic error detection and diagnosis of file servers.

15

16 Figure 2 shows a process flow diagram of a first method for operating a
17 system for computer assisted automatic error detection and diagnosis of file servers, in-
18 cluding periodic and continuous interpretation of monitoring statistics.

19

20 Figure 3 shows a process flow diagram of a second method for operating a
21 system for computer assisted automatic error detection and diagnosis of file servers, in-
22 cluding augmentation of known network protocols.

1
2 Figure 4 shows a process flow diagram of a third method for operating a
3 system for computer assisted automatic error detection and diagnosis of file servers, in-
4 cluding cross-layer analysis of monitoring statistics.
5

6 Figure 5 shows a process flow diagram of a fourth method for operating a
7 system for computer assisted automatic error detection and diagnosis of file servers, in-
8 cluding tracking of hardware and software configuration changes.
9
10

Detailed Description of the Preferred Embodiment

11
12 In the following description, a preferred embodiment of the invention is de-
13 scribed with regard to preferred process steps and data structures. Embodiments of the
14 invention can be implemented using general-purpose processors or special purpose proc-
15 essors operating under program control, or other circuits, adapted to particular process
16 steps and data structures described herein. Implementation of the process steps and data
17 structures described herein would not require undue experimentation or further invention.
18

1 *Lexicography*

2

3 The following terms refer or relate to aspects of the invention as described
4 below. The descriptions of general meanings of these terms are not intended to be limit-
5 ing, only illustrative.

6

7 • **client** and **server** — in general, these terms refer to a relationship between two de-
8 vices, particularly to their relationship as client and server, not necessarily to any
9 particular physical devices.10
11 For example, but without limitation, a particular client device in a first relationship
12 with a first server device, can serve as a server device in a second relationship with
13 a second client device. In a preferred embodiment, there are generally a relatively
14 small number of server devices servicing a relatively larger number of client de-
15 vices.

16

17 • **client device** and **server device** — in general, these terms refer to devices taking
18 on the role of a client device or a server device in a client-server relationship (such
19 as an HTTP web client and web server). There is no particular requirement that
20 any client devices or server devices must be individual physical devices. They can
21 each be a single device, a set of cooperating devices, a portion of a device, or some
22 combination thereof.

1
2 For example, but without limitation, the client device and the server device in a
3 client-server relation can actually be the same physical device, with a first set of
4 software elements serving to perform client functions and a second set of software
5 elements serving to perform server functions

6
7 • **configuration changes** — in general, information regarding changes to a configu-
8 ration of the file server

9
10 • **cross-layer analysis** — in general, a technique for applying a large number of
11 combinations of diagnostic constraints, so as to determine a set of errors or other
12 faults are not compatible with the set of current monitoring statistics

13
14 • **diagnostic constraint** — in general, a functional or logical constraint on possible
15 errors or other failures, due to functional or logical structure of the file server or its
16 operating system

17
18 • **diagnostic software module** — in general, a software module in the file server
19 used for performing computer assisted automatic error detection and diagnosis

20
21 • **error detection and diagnosis** — in general, a technique for detecting errors and
22 other failures, and for determining a likely cause thereof

1

2 • **lower-level** and **higher-level** protocols — in general, these terms refer to a rela-
3 tionship between two protocols, particularly to their relationship as a higher-level
4 protocol which relies on operation of a lower-level protocol and which is able to
5 alter parameters of the lower-level protocol, not necessarily to any particular pro-
6 tocols.

7

8 • **manipulating parameters** — in general, a technique for using a higher-level
9 protocol to determine whether a lower-level protocol is operating relatively effi-
10 ciently using a set of selected parameters for the lower-level protocol, and using
11 the lower-level protocol to repeatedly and rapidly alter those selected parameters
12 so as to find an optimal set of selected parameters

13

14 • **monitoring statistics** — in general, information regarding performance of the file
15 server or other device

16

17 • **network protocol** — in general, a technique for communication between devices,
18 such as for example between (a) the file server or other device and (b) a point ex-
19 ternal to the file server or other device.

20

21 • **pattern matching** — in general, a technique for comparing a set of monitoring
22 statistics against a selected pattern known to be related to an error or other fault

1

2 • **periodic and continuous interpretation** — in general, repeated and rapid review

3 of monitoring statistics, so as to identify errors or other faults rapidly and as early

4 as possible

5

6 • **protocol augmentation** — in general, a technique for using a higher-level proto-

7 col to determine whether a lower-level protocol is operating relatively efficiently

8 using a set of selected parameters for the lower-level protocol, and using the

9 lower-level protocol to repeatedly and rapidly alter those selected parameters so as

10 to find an optimal set of selected parameters

11

12 • **rule-based inference system** — in general, a technique for drawing factual con-

13 clusions from monitoring statistics and other known facts, so as to determine the

14 presence or absence of an error or other fault

15

16 • **sequential and combined constraints on diagnosis** — in general, a technique for

17 applying a large number of combinations of diagnostic constraints, so as to deter-

18 mine a set of errors or other faults are not compatible with the set of current

19 monitoring statistics

20

21 • **usage profile** — in general, information regarding a pattern or profile of use of a

22 file server or other device

1
2 As noted above, these descriptions of general meanings of these terms are
3 not intended to be limiting, only illustrative. Other and further applications of the inven-
4 tion, including extensions of these terms and concepts, would be clear to those of ordinary
5 skill in the art after perusing this application. These other and further applications are
6 part of the scope and spirit of the invention, and would be clear to those of ordinary skill
7 in the art, without further invention or undue experimentation.

8
9 *System Elements*
10
11
12

13
14 Figure 1 shows a block diagram of a system for computer assisted auto-
15 matic error detection and diagnosis of file servers.

16
17 A system 100 includes a file server (or other device) 110, a communication
18 network 120, a network interface 130, and a database 140.

19
20 The file server (or other device) 110 includes a computer having a proces-
21 sor, program and data memory, mass storage, a presentation element, and an input ele-
22 ment, and coupled to the communication network 120. As used herein, the term "com-
puter" is intended in its broadest sense, and includes any device having a programmable
processor or otherwise falling within the generalized Turing machine paradigm. The
mass storage can include any device for storing relatively large amounts of information,

1 such as magnetic disks or tapes, optical devices, magneto-optical devices, or other types
2 of mass storage.

3
4 The file server 110 includes operating system software 111, including an
5 appliance operating system 112 for controlling the file server 110 and performing its file
6 server operations, and a diagnostic software module 113 for performing computer assisted
7 automatic error detection and diagnosis.

8
9 The diagnostic software module 113 includes software modules for error
10 detection and diagnosis (further described with regard to figure 2, figure 3, figure 4, and
11 figure 5), and is coupled to the database 140 for storing and retrieving statistical and other
12 information.

13
14 The communication network 120 includes any technique for sending infor-
15 mation between the file server 110 and at least one point outside the file server 110. In a
16 preferred embodiment, the communication network 120 includes a computer network,
17 such as an Internet, intranet, extranet, or a virtual private network. In alternative em-
18 bodiments, the communication network 120 can include a direct communication line, a
19 switched network such as a telephone network, or some combination thereof.

20
21 The network interface 130 includes a communication link between the file
22 server 110 and the communication network 120.

1
2 The database 140 includes memory or mass storage in which monitoring
3 statistics and other information about the file server 110 can be recorded and retrieved
4 therefrom.

5
6 *Method of Operation (Interpretation of Monitoring Statistics)*

7
8 Figure 2 shows a process flow diagram of a first method for operating a
9 system for computer assisted automatic error detection and diagnosis of file servers, in-
10 cluding periodic and continuous interpretation of monitoring statistics.

11
12 A method 200 includes a set of flow points and a set of steps. The system
13 100 performs the method 200. Although the method 200 is described serially, the steps of
14 the method 200 can be performed by separate elements in conjunction or in parallel,
15 whether asynchronously, in a pipelined manner, or otherwise. There is no particular re-
16 quirement that the method 200 be performed in the same order in which this description
17 lists the steps, except where so indicated.

18
19 At a flow point 210, the diagnostic software module 113 is ready to pas-
20 sively monitor operating system statistics.

21

1 At a step 211, the diagnostic software module 113 receives operating sys-
2 tem statistics from the database 140 and compares those operating system statistics
3 against a set of predefined rules regarding statistical aberrations. If the operating system
4 statistics fall within the predefined rules, thus indicating an abnormal system state, the di-
5 agnostic software module 113 flags that abnormal system state, enters a reporting record
6 in the database 140, and reports the detected abnormal system state and its associated op-
7 erating system statistics to an operator (not shown).

8
9 In a preferred embodiment, the predefined rules regarding statistical aberra-
10 tions include both (a) pattern matching, in which the diagnostic software module 113
11 compares the monitoring statistics against selected patterns known to be related to an er-
12 rors or other faults, and (b) rule-based inference, in which the diagnostic software module
13 113 draws factual conclusions from monitoring statistics and other known facts, so as to
14 determine the presence or absence of errors or other faults.

15
16 In a preferred embodiment, the predefined rules regarding statistical aberra-
17 tions are responsive to a usage pattern for the file server, as recorded in the database 140
18 or otherwise. For example, the file server can be primarily used as for an ISP, a devel-
19 opment environment, a mail server, or otherwise. The usage pattern for the file server
20 could reflect such information, preferably including an expected number of distinct users,
21 an expected number of distinct files, an expected size for typical files, and an expected
22 duration of existence for distinct files or file-usage sessions.

1
2 At a step 212, the diagnostic software module 113 determines, using the
3 predefined rules, whether the abnormal system state corresponds to a specific known
4 problem, or is associated with more than one known problem. If the abnormal system
5 state corresponds to a specific known problem, the diagnostic software module 113 so in-
6 dicates and reports the specific known problem to the operator.

7
8 At a step 213, if the abnormal system state was determined in the previous
9 step to correspond to a set of more than one known problem, the diagnostic software
10 module 113 determines if there is a predefined technique for analyzing the abnormal sys-
11 tem state. If there is such a predefined technique, the diagnostic software module 113
12 proceeds with the next step.

13
14 At a step 214, the diagnostic software module 113 performs the predefined
15 technique for analyzing the abnormal system state found in the previous step.

16
17 After this step, the method 200 has performed one round of passively
18 monitoring operating system statistics, and is ready to perform another such round so as
19 to continuously monitor operating system statistics.

20
21 The method 200 is performed one or more times starting from the flow
22 point 210 and continuing therefrom. In a preferred embodiment, the diagnostic software

1 module 113 repeatedly performs the method 200, starting from the flow point 210 and
2 continuing therefrom, so as to review monitoring statistics gathered by the file server pe-
3 riодically and continuously.

4

5 *Method of Operation (Augmentation of Network Protocols)*

6

7 Figure 3 shows a process flow diagram of a second method for operating a
8 system for computer assisted automatic error detection and diagnosis of file servers, in-
9 cluding augmentation of known network protocols.

10
11
12
13
14
15
16

17 A method 300 includes a set of flow points and a set of steps. The system
18 100 performs the method 300. Although the method 300 is described serially, the steps of
19 the method 300 can be performed by separate elements in conjunction or in parallel,
20 whether asynchronously, in a pipelined manner, or otherwise. There is no particular re-
21 quirement that the method 300 be performed in the same order in which this description
lists the steps, except where so indicated.

17

18 At a flow point 310, the diagnostic software module 113 is ready to deter-
19 mine the presence or absence of errors or other faults, using protocol augmentation of
20 known network protocols.

21

1 At a step 311, the diagnostic software module 113 selects parameters for
2 known network protocols, so as to communicate with a device at a point outside the file
3 server 110. In selecting these parameters, the diagnostic software module 113 attempts to
4 choose parameters that match with the device outside the file server 110, so as to achieve
5 relatively efficient communication.

6

7 At a step 312, the diagnostic software module 113 uses the lower-level
8 protocol to set its parameters to the parameters selected in the previous step.

9

10 At a step 313, the diagnostic software module 113 communicates with the
11 device outside the file server 110 using a higher-level protocol, and uses the higher-level
12 protocol to measure the relative efficiency of communication therewith.

13

14 At a step 314, the diagnostic software module 113 selects a new set of pa-
15 rameters for the lower-level protocol, in response to the measure of relative efficiency
16 determined using the higher-level protocol in the previous step.

17

18 The diagnostic software module 113 performs the step 312, the step 313,
19 and the step 314, repeatedly and rapidly, so as to try a large number of combinations of
20 protocol parameters. The diagnostic software module 113 thus determines which selected
21 parameters are optimal.

22

1 At a step 315, the diagnostic software module 113 uses the lower-level
2 protocol to set its parameters to the parameters determined to be optimal.

3

4 After this step, the method 300 has set the optimal parameters for the lower-
5 level protocol.

6

7 The method 300 can be performed one or more times starting from the flow
8 point 310 and continuing therefrom.

9

10 *Method of Operation (Cross-Layer Analysis)*

11

12 Figure 4 shows a process flow diagram of a third method for operating a
13 system for computer assisted automatic error detection and diagnosis of file servers, in-
14 cluding cross-layer analysis of monitoring statistics.

15

16 A method 400 includes a set of flow points and a set of steps. The system
17 100 performs the method 400. Although the method 400 is described serially, the steps of
18 the method 400 can be performed by separate elements in conjunction or in parallel,
19 whether asynchronously, in a pipelined manner, or otherwise. There is no particular re-
20 quirement that the method 400 be performed in the same order in which this description
21 lists the steps, except where so indicated.

22

1 At a flow point 410, the diagnostic software module 113 is ready to perform
2 cross-layer analysis of monitoring statistics.

3
4 At a step 411, the diagnostic software module 113 identifies a diagnostic
5 constraint applicable to the set of current monitoring statistics.

6
7 At a step 412, the diagnostic software module 113 performs pattern match-
8 ing and rule-based inference to attempt to identify a known error or other fault, similar to
9 the method 200.

10
11 The diagnostic software module 113 performs the step 411 and the step 412
12 repeatedly and rapidly, so as to apply a large number of combinations of diagnostic con-
13 straints. The diagnostic software module 113 thus determines a set of errors or other
14 faults are not compatible with the set of current monitoring statistics.

15
16 In a preferred embodiment, the diagnostic software module 113 continues
17 to apply combinations of diagnostic constraints until the set of errors or other faults is re-
18 duced to only a few possibilities.

19
20 After this step, the method 400 has used cross-layer analysis of monitoring
21 statistics to restrict the set of possible errors or other faults.

22

1 The method 400 can be performed one or more times starting from the flow
2 point 410 and continuing therefrom.

3

4 *Method of Operation (Tracking Configuration Changes)*

5

6 Figure 5 shows a process flow diagram of a fourth method for operating a
7 system for computer assisted automatic error detection and diagnosis of file servers, in-
8 cluding tracking of hardware and software configuration changes.

9

10 A method 500 includes a set of flow points and a set of steps. The system
11 100 performs the method 500. Although the method 500 is described serially, the steps of
12 the method 500 can be performed by separate elements in conjunction or in parallel,
13 whether asynchronously, in a pipelined manner, or otherwise. There is no particular re-
14 quirement that the method 500 be performed in the same order in which this description
15 lists the steps, except where so indicated.

16

17 At a flow point 510, the diagnostic software module 113 is ready to perform
18 tracking of hardware and software configuration changes.

19

20 At a step 511, the diagnostic software module 113 identifies a past hard-
21 ware or software configuration for the file server, recorded in the database 140 or other-

1 wise, and a present hardware or software configuration for the file server, recorded in the
2 database 140 or otherwise.

3

4 At a step 512, the diagnostic software module 113 identifies a set of differ-
5 ences between the past configuration and the present configuration.

6

7 At a step 513, the diagnostic software module 113 identifies a set of possi-
8 ble errors or other faults associated with the set of differences identified in the previous
9 step.

10

11 At a step 514, the diagnostic software module 113 applies diagnostic con-
12 straints from the set of current monitoring statistics to the set of possible errors or other
13 faults, so as to determine which errors or other faults are most likely.

14

15

16

17

18

19

20

21

22

At a step 515, the diagnostic software module 113 uses the determination,
in the previous step, of which errors or other faults are most likely, to suggest activities to
reverse the hardware or software configuration changes so as to place the file server in an
operating state.

After this step, the method 500 has used tracking of hardware and software
configuration changes to restrict the set of possible errors or other faults.

1 The method 500 can be performed one or more times starting from the flow
2 point 510 and continuing therefrom.

3

4 *Generality of the Invention*

5

6 The invention has general applicability to various fields of use, not neces-
7 sarily related to the services described above. For example, these fields of use can in-
8 clude other “appliance”-type devices other than file servers.

9
10
11
12
13
14
15
16
17
18
19

10 Other and further applications of the invention in its most general form, will
11 be clear to those skilled in the art after perusal of this application, and are within the
12 scope and spirit of the invention.

13
14 *Technical Appendix*

15
16 The technical appendix enclosed with this application is hereby incorpo-
17 rated by reference as if fully set forth herein, and forms a part of the disclosure of the in-
18 vention and its preferred embodiments.

1 Alternative Embodiments

2

3 Although preferred embodiments are disclosed herein, many variations are
4 possible which remain within the concept, scope, and spirit of the invention, and these
5 variations would become clear to those skilled in the art after perusal of this application.